

(43) Date of A Publication 10.01.2001

(21) Application No 0005570.7

(22) Date of Filing 08.03.2000

(30) Priority Data

(31) 09283598

(32) 01.04.1999

(33) US

(71) Applicant(s)

NDS Limited
(Incorporated in the United Kingdom)
1 Heathrow Boulevard, 286 Bath Road,
WEST DRAYTON, Middlesex, UB7 0DQ,
United Kingdom

(72) Inventor(s)

Steve Epstein
Shlomo Kipnis

(74) Agent and/or Address for Service

Reginald W Barker & Co
Clifford's Inn, Fetter Lane, LONDON, EC4A 1BX,
United Kingdom

(51) INT CL⁷

H04H 1/00, H04N 7/173

(52) UK CL (Edition S)

H4R RCSS R17D R17E

(56) Documents Cited

WO 98/37696 A WO 98/28869 A WO 97/17774 A
US 5758257 A

(58) Field of Search

UK CL (Edition R) H4F FBB, H4R RCSS RCST RCT
RCX RLS
INT CL⁷ G06F 17/60, H04H 1/00 1/02, H04N 7/173
online: EPODOC, WPI, JAPIO

(54) Abstract Title

Item selection for broadcasting system

(57) A method for selecting an item to be broadcast in a broadcasting system includes : assigning 175 a communication channel having a bandwidth to a community of users of similar interests; for each one item of a plurality of available items, determining 180 an affinity between the one item and the community and associating 185 the affinity with the one item; computing 190 a measure of remaining bandwidth of the communication channel assigned to the community; and choosing 195 at least one item to be broadcast from among the plurality of available items based on the affinity associated with the one item and on the measure of remaining bandwidth. The communities are established from profiles of individual users. The method is particularly suitable where the users can receive only a single channel at a time.

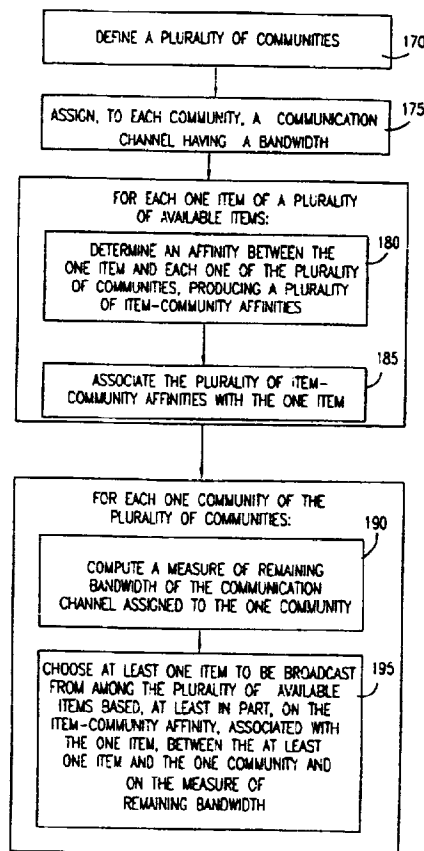


FIG. 3





FIG. 1



3/10

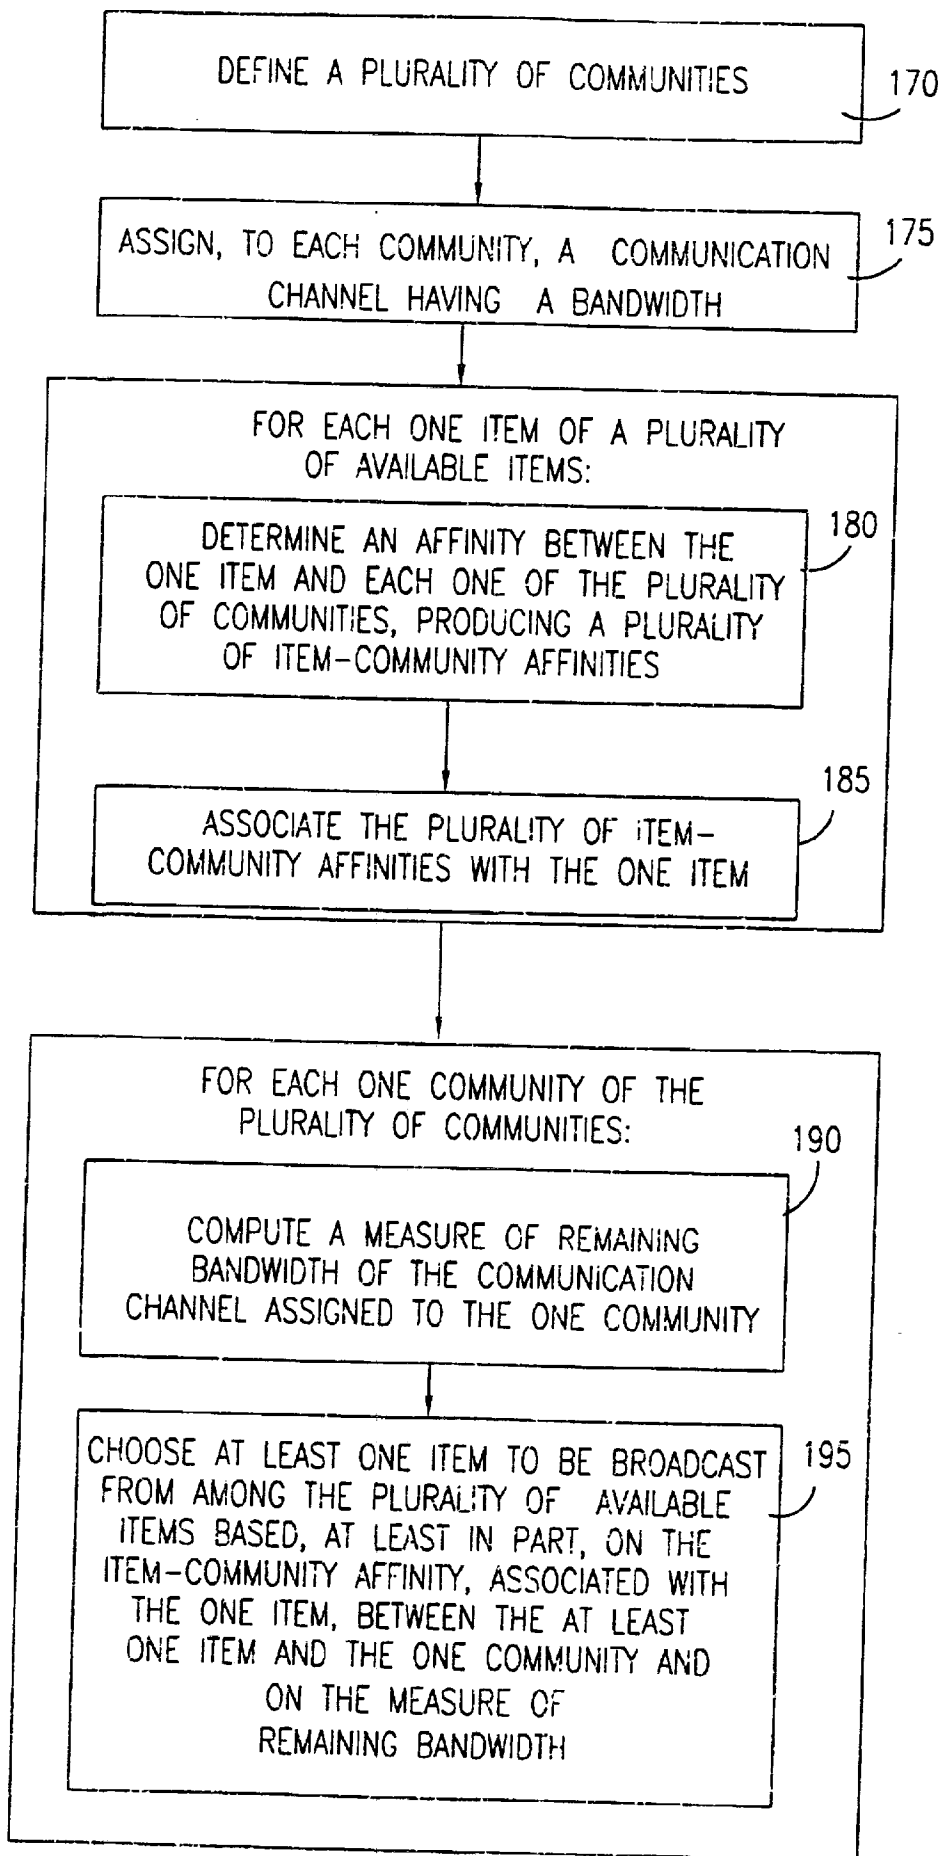


FIG. 3

FIG. 4

TO BE PERFORMED FOR EACH ITEM:

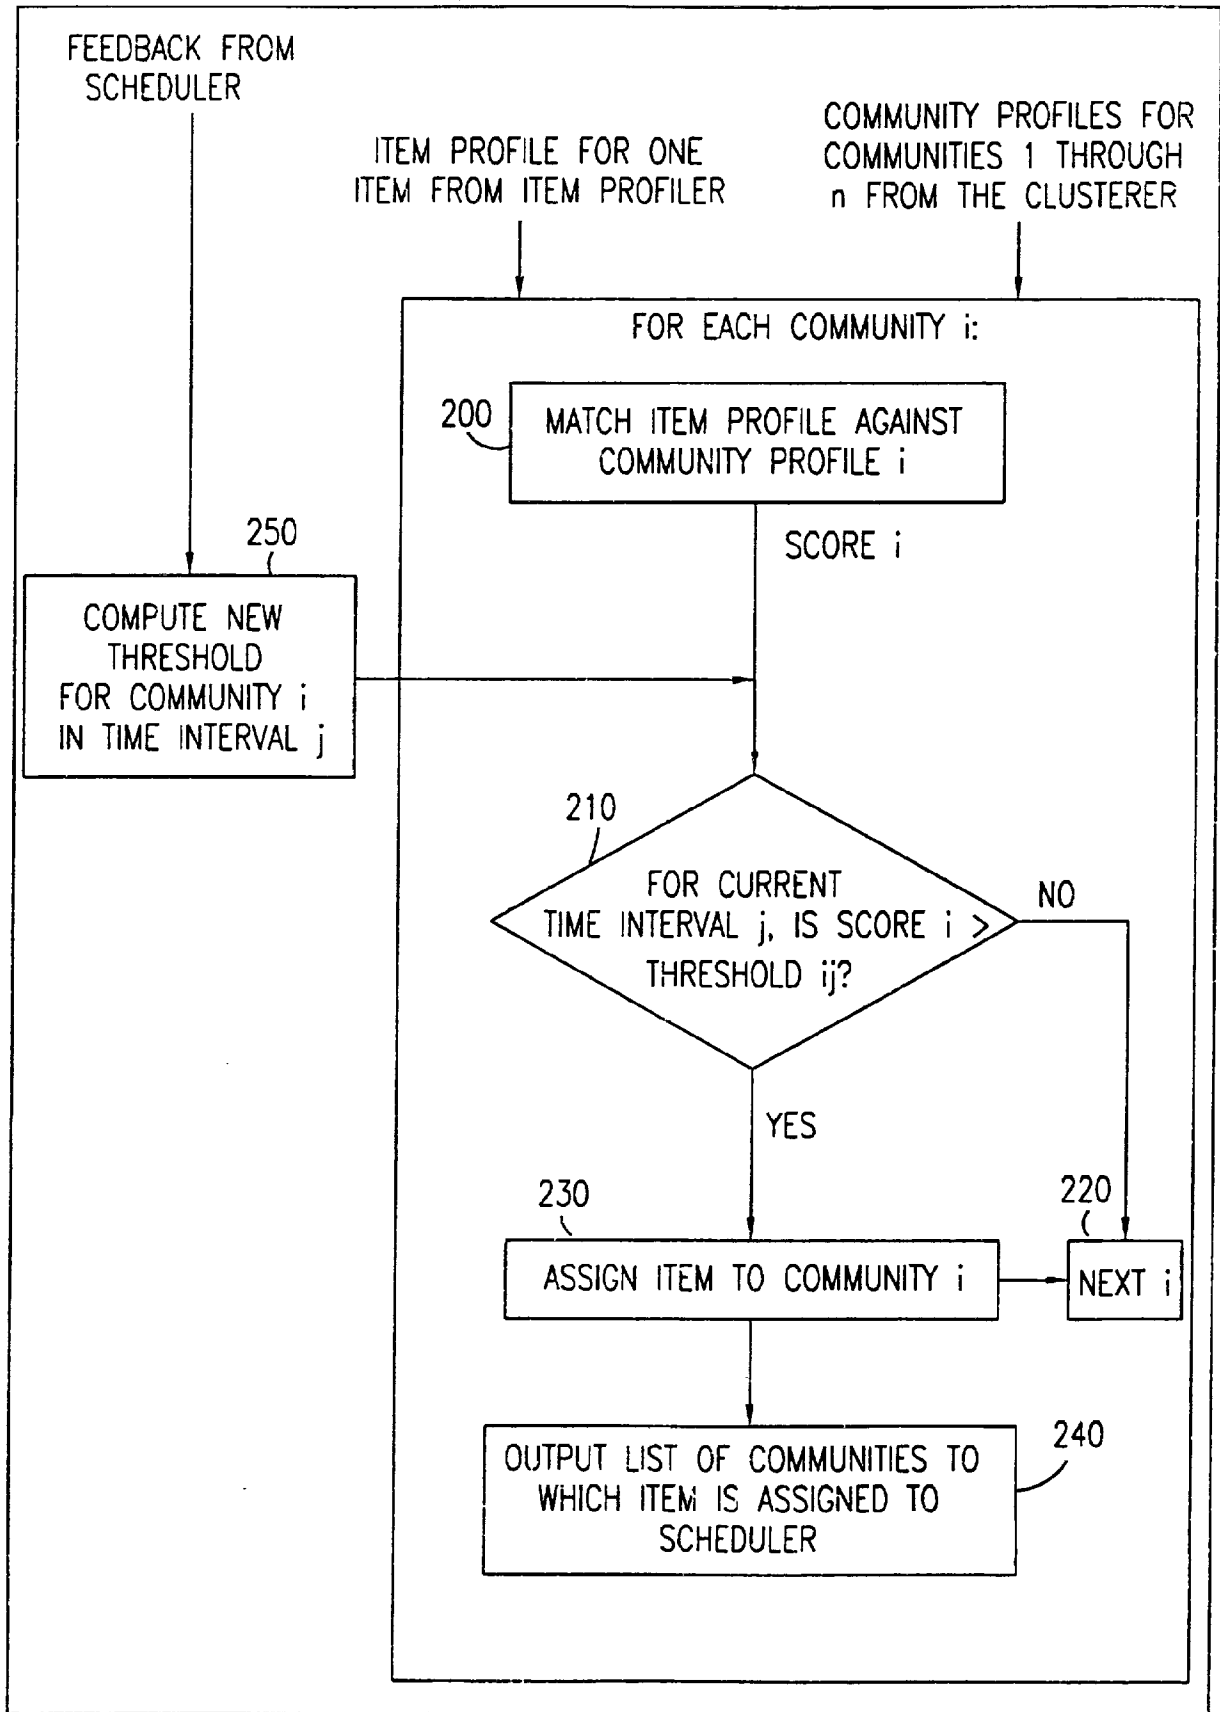


FIG. 5

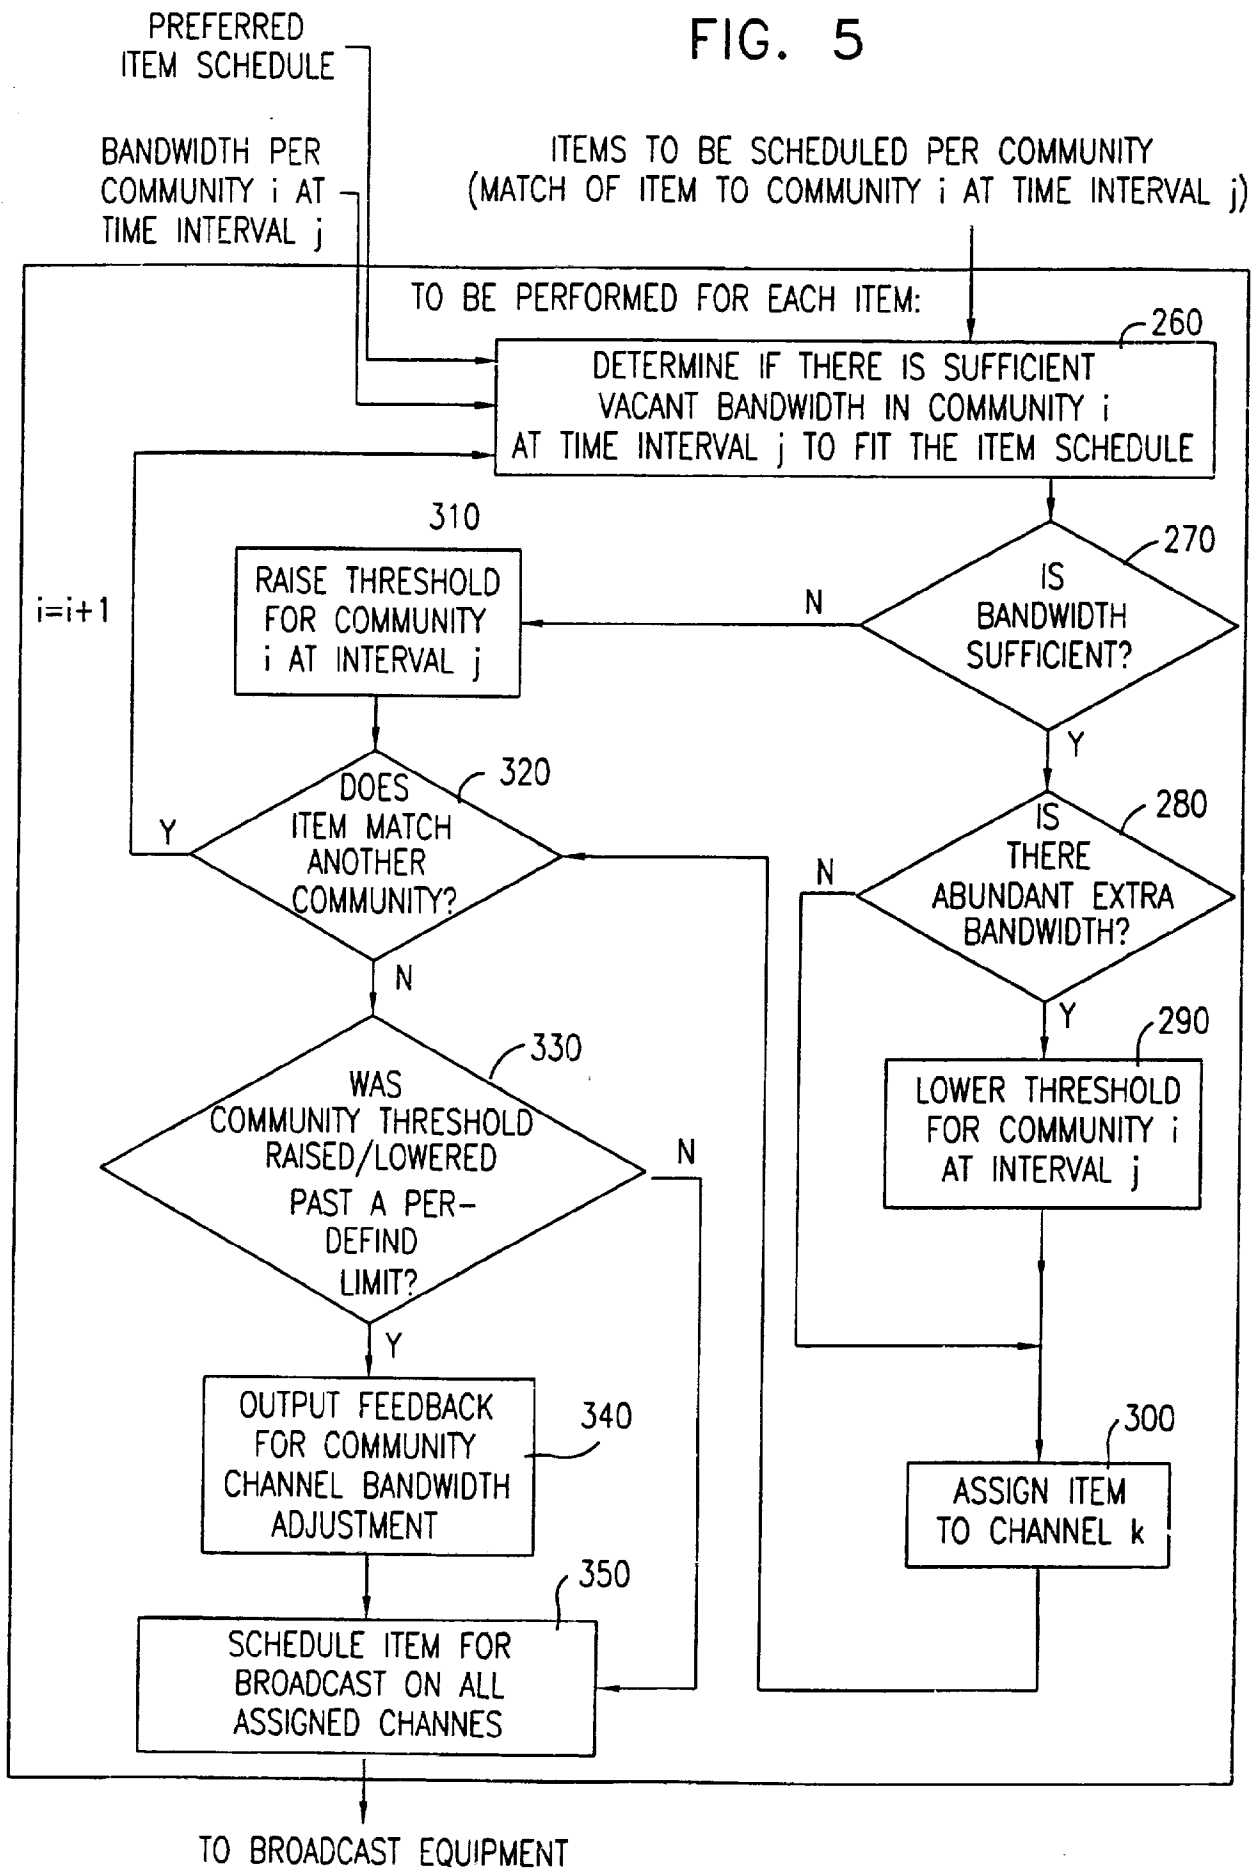


FIG. 7

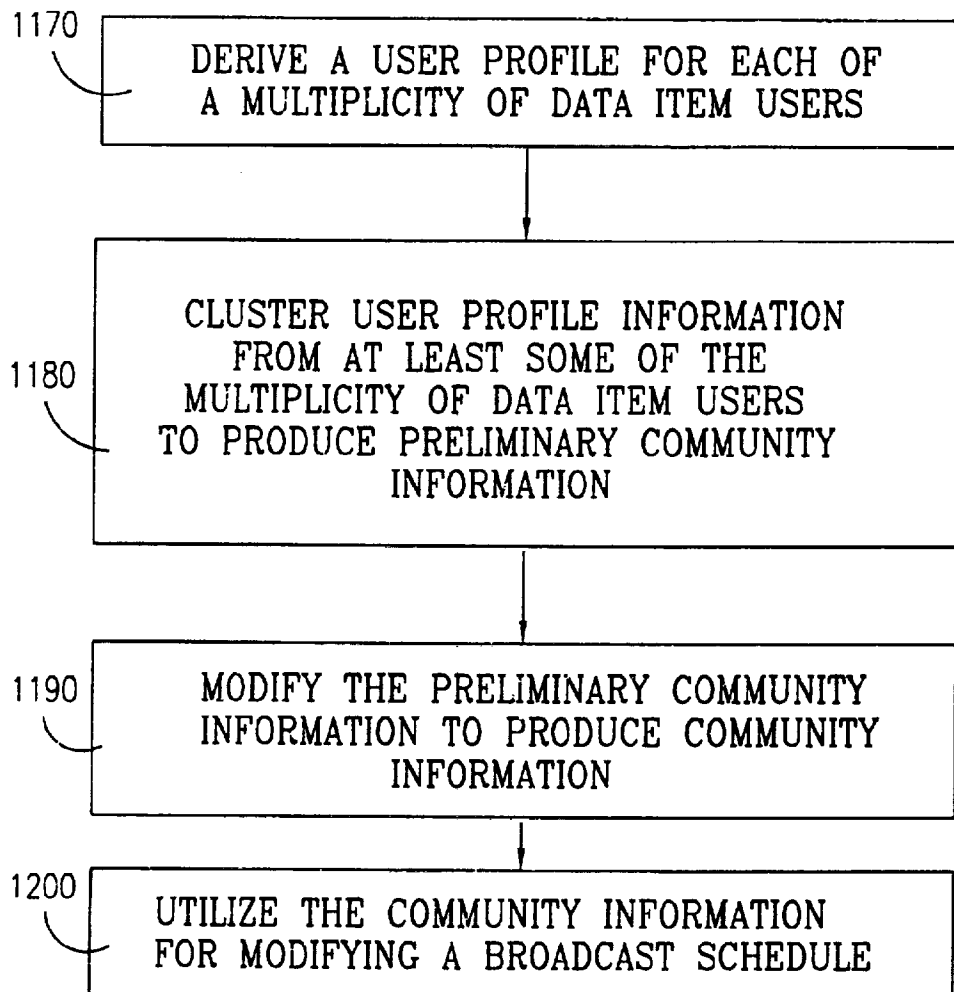
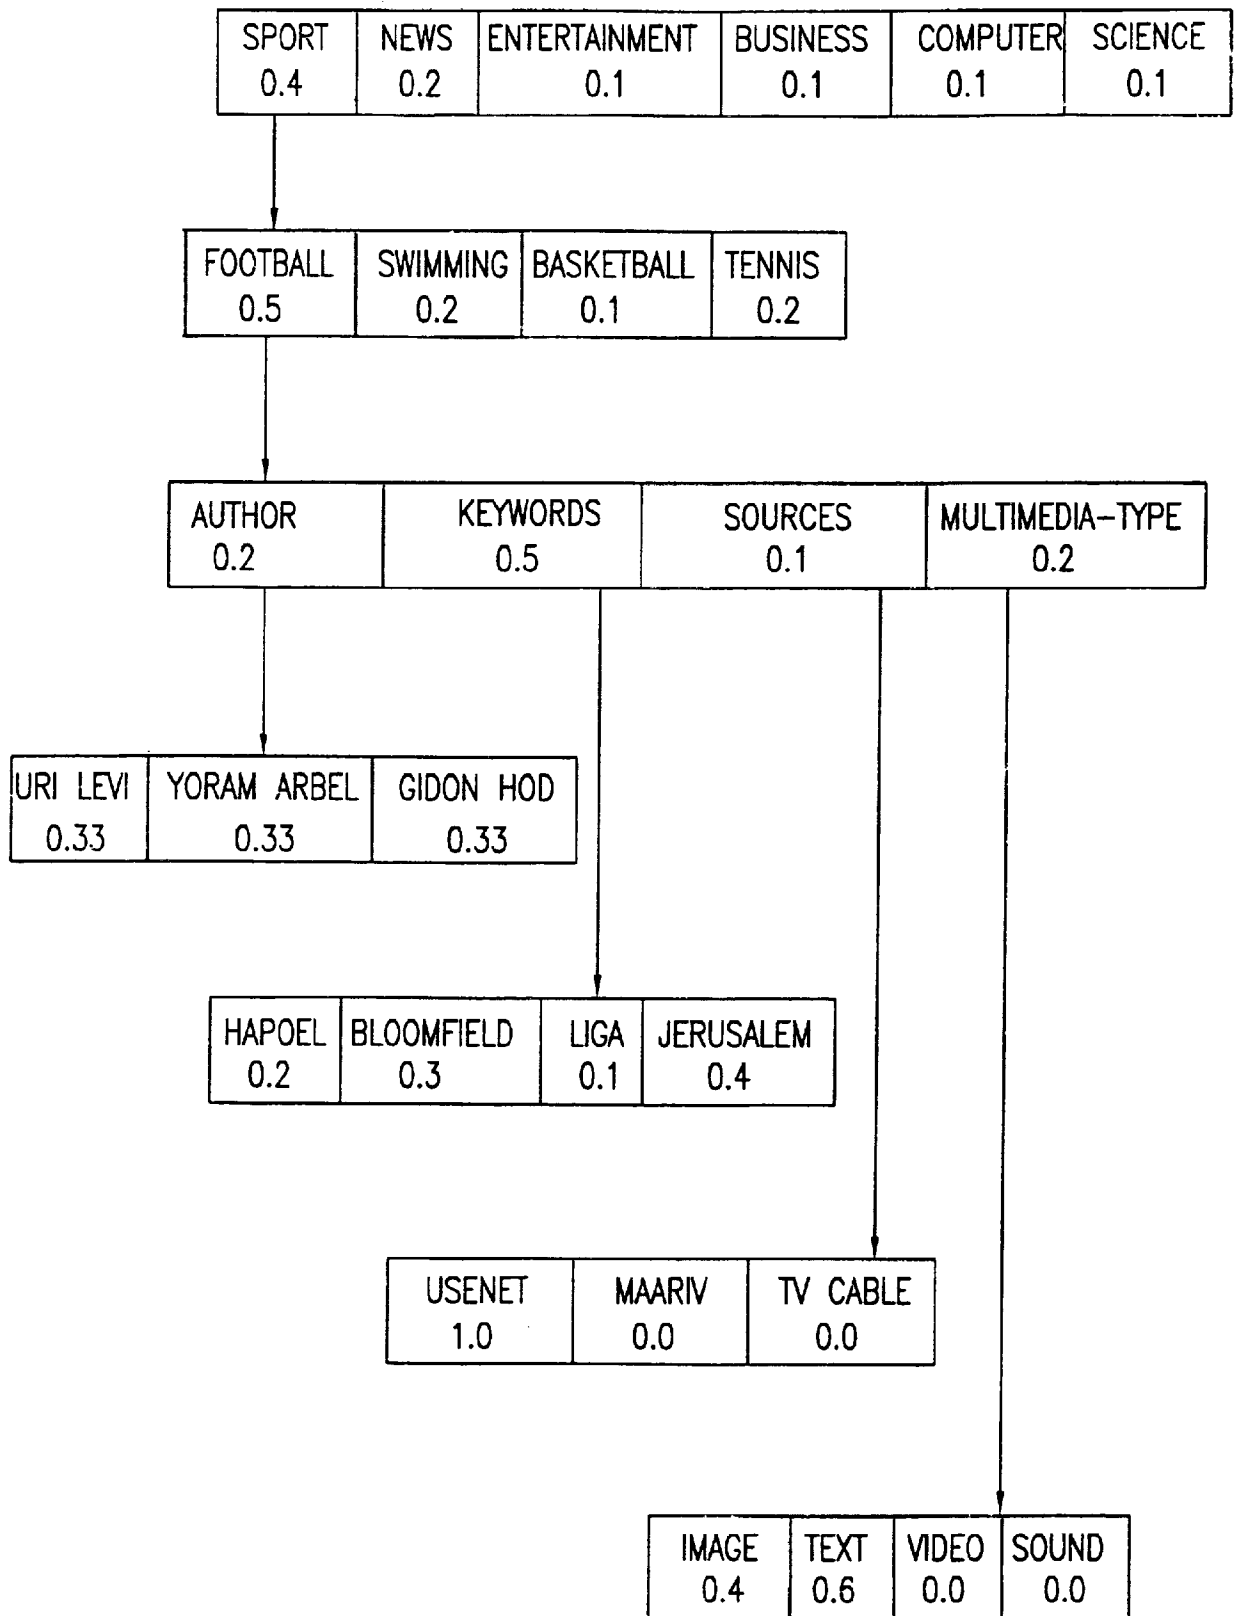
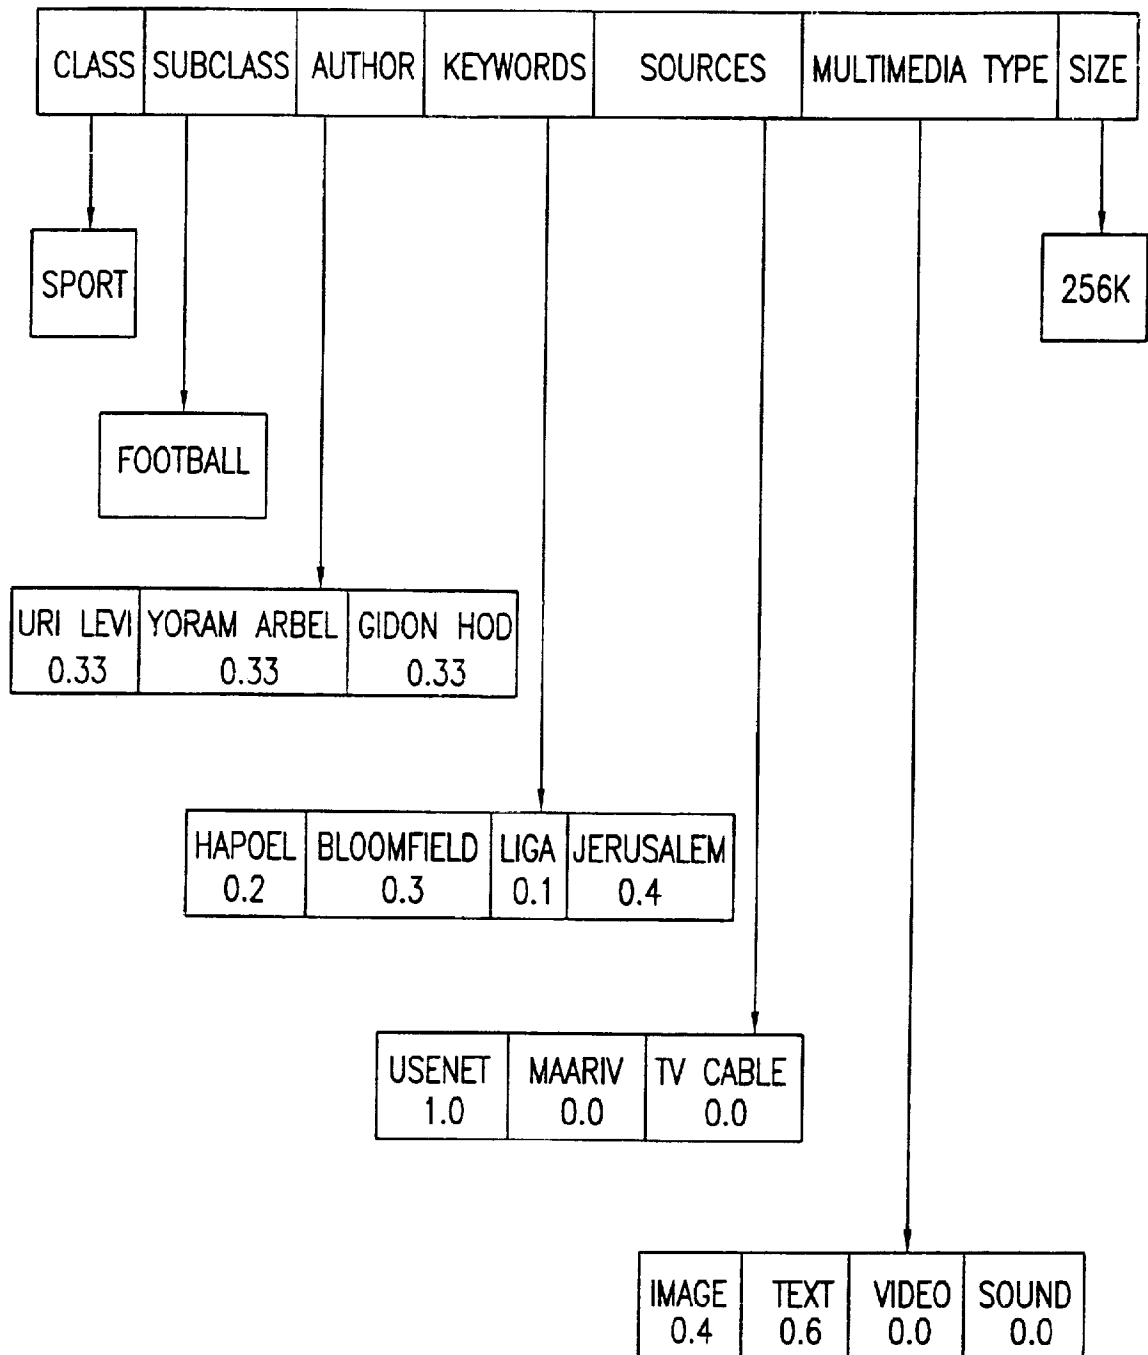


FIG. 8



9/10

FIG. 9



10/10

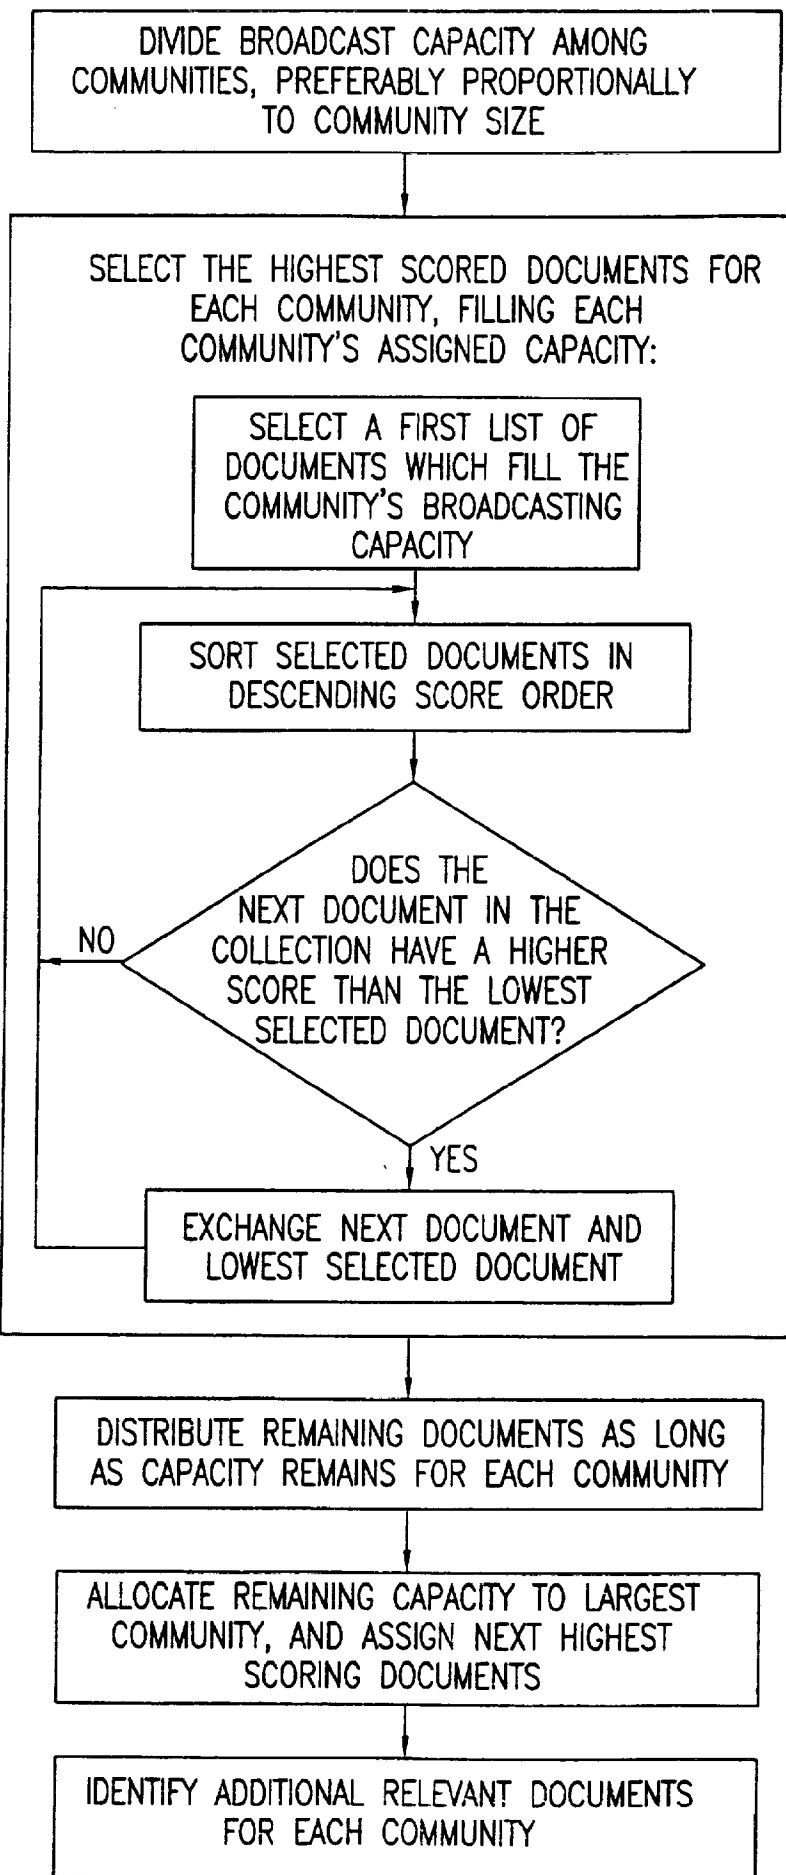


FIG. 10

FIELD OF THE INVENTION

The present invention relates to broadcasting systems in general, and in particular to programming, scheduling, and bandwidth allocation systems, particularly systems suitable for data broadcasting.

5

BACKGROUND OF THE INVENTION

Broadcasting systems, including data broadcasting systems, are well known in the art. Examples of such systems include systems available under the following names: Cyberstar, Globecast, AstraNet, DirectPC, DSSPC, and Media4.

10

It is well-known in broadcasting systems to use a variety of methods for determining what items to broadcast and at what time the items should be broadcast. Generally, broadcasting systems are characterized by having a given available broadcast bandwidth which must be allocated between physical and/or logical broadcast channels. The systems are designed for broadcasting to users or customers who typically have different preferences for broadcasting content to be received. In the prior art work has been directed towards increasing bandwidth utilization, creating logical channels of interest, and increasing customer satisfaction. In general, it is desirable to determine what to broadcast, when to broadcast, and on which channel to broadcast, in order to increase utilization of available bandwidth while reaching high levels of user/customer satisfaction from content received.

15

Typically, in state of the art broadcasting systems, information is collected periodically from users/customers or from a representative sample thereof. Broadcast programs and schedules are then designed, based in part of the collected information, to increase bandwidth utilization and user/customer satisfaction. Typically, the process of arranging programming and scheduling is time consuming and is semi-manual; the process is typically based on only a representative sample, as described above. The state of the art methods described are typically sufficient for television broadcasting systems, in which each of a small number of channels carries a small number of items (programs), each of relatively long time duration. In data broadcasting systems, on the other hand, where there are typically a large number of channels and a large number of short-

20

25

30

duration items, such semi-manual methods are known to be unsatisfactory or even impossible.

One system for scheduling broadcasts using customer profiles is described in US Patent 5,758,257 to Herz et al. The Herz et al patent describes scheduling the receipt of desired movies or other forms of data by means of individual customer profiles describing each individual customer. A so-called “agreement matrix” is calculated by comparing the recipient’s profiles to the actual profiles of the available programs or other data. A virtual channel for each individual is generated from the “agreement matrix”, in an attempt to satisfy the desires of each individual via their own virtual channel.

A typical example of conditional access methods generally applicable to broadcasting is described in US Patents 5,282,249 and 5,481,609 to Cohen et al.

An example of conditional access techniques particularly suited to a data broadcasting context is described in a co-pending US Patent Application assigned to the assignee of the present application, submitted on February 4, 1999 and entitled “Key Delivery in a Secure Broadcasting System”, with inventors Yishai Sered and Nigel Stephen Dallard; the patent application is termed herein the “Key Delivery application”. Corresponding applications include published PCT application WO 99/45711 and Israel patent application 123554.

Related technologies are described in a document numbered IN-T215 and entitled “Functional Design Overview for Content Providers and Application Developers”, published by NDS Ltd. on November 3, 1998.

The following references provide a sample of the state of the art, and are useful in understanding the present invention:

1. S. Acharya, M. Franklin and S. Zdonik. 1996. Disseminating updates on broadcast disks. *Proceedings of the 22nd VLDB Conference*, Mumbai (Bombay), India.

2. S. Acharya, R. Alonso, M. J. Franklin, and S. B. Zdonik. 1995. Broadcast disks: Data management for asymmetric communications environments. In M. J. Carey and D. A. Schneider, editors, *Proceedings of the 1995 ACM*

SIGMOD International Conference on Management of Data, pages 199-210, San Jose, California, 1995. SIGMOD Record 24(2).

3. W. P. Birmingham, E. H. Durfee, T. Mullen, and M. P. Wellman. 1995. The distributed agent architecture of the University of Michigan Digital Library (UMDL). In *Spring Symposium Series on Software Agent*.

4. E. Bloedorn, I. Mani and T. R. MacMillan. 1996. Representational issues in machine learning of user profiles. *AAAI Spring Symposium on Machine Learning in Information Access*.

5. E. W. Brown, J. P. Callan, B. Croft. 1994. Fast incremental indexing for full-text information retrieval. Proceedings of *VLDB Conf*, pages 192-202.

6. cc:Browser. 1997. <http://www.faulkner.com/>

7. S. D. Cohen and L. J. Guibas. 1996. Shape-based illustration indexing and retrieval some first steps Stanford University, Stanford, *ARPA Image Understanding Workshop 1996*.

8. CNN++ News filtering Agent. <http://www.cs.umbc.edu/~smart-ht/info.html>

9. D. Cutting and J. Pederson. 1990. Optimizations for dynamic inverted index maintenance. Proceedings of *SIGIR*, pages 405-411.

10. Y. Ejgenberg and Y. Lindel. 1997. B.S.c project, Computer Science Department at Bar Ilan University.

11. P. Edwards, D. Bayer, C. L. Green and T. R. Payne. 1996. Experience with learning agents which manage Internet-based information. *AAAI Spring Symposium on Machine learning in IA*, Scotland.

12. C. Faloutsos and D. W. Oard. 1995. A survey of information retrieval and filtering methods. *University of Maryland CS-TR-3514*.

13. W. B. Frakes and R. Baeza-Yates. 1992. *Information Retrieval: Data Structures & Algorithms*. Prentice Hall Englewood Cliffs, New Jersey.

14. K. C. Gowda and G. Krishna. 1978. Agglomerative clustering using the concept of mutual nearest neighbourhood. *Pattern Recognition*, Vol. 10, No 2, pp. 105-112

15. S. Gant. WEBFilter: An agent for filtering the Web.
<http://ils.unc.edu/webfilter/>
16. D. E. Goldberg. 1989. *Genetic Algorithm in Search, Optimization and Machine Learning*. Addison-Wesley.
- 5 17. T. Imielinski , S. Viswanathan. 1994. Adaptive wireless information systems. In Proceedings of GIGDBS Conference, Tokyo, Japan.
18. I. Ink and N. McGough. 1992. Robots and mail filtering
<http://www.ii.com/internet/robots/>
19. A. Kobsa. 1994. User modeling and user-adapted interaction. In
10 *Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems*, Volume 2 of TUTORIALS, pages 415-416.
20. K. Lang. 1995. News Weeder: Learning to filter netnews. *Proceedings of INT Conference of Machine Learning*, PP. 331-339.
21. W. Y. Ma and B. S. Manjunath. 1995. Image indexing using a
15 texture dictionary. *Proceedings of SPIE Conference on Image Storage and Archiving System* ,Vol. 2606, pp. 288-298. Philadelphia Pennsylvania.
22. P. Maes and R. Kozierok. 1993. Learning interface agents. *Proceedings of AAAI-93 Washington D.C.* pp. 459-465.
23. B. S. Manjunath and W. Y. Ma. 1996. Texture features for
20 browsing and retrieval of image data. *IEEE T-PAMI special issue on Digital Libraries*.
24. A. Moukas. 1996. Amalthaea: Information discovery and filtering using a multi-agent evolving ecosystem. *The first international conference on the Practical Application of Intelligent Agents and Multi Agents Technology*.
25 pp. 421-436.
25. S. J. Pelletier and J. F. Arcand. 1996. STEALTH: A personal digital assistant for information filtering. *The first international conference on the Practical Application of Intelligent Agents and Multi Agents Technology*. pp. 455-474.
- 30 26. E. Rasmussen. 1992. *Information Retrieval. Data Structures and Algorithms*. Editors: W. B. Frakes and R. Baeza-Yates Prentice Hall Inc., Englewood Cliffs, N. J.

27. P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom and J. Riedl. 1994. GroupLens: An open architecture for collaborative filtering of netnews. *Proceeding of ACM CSCW 94 Conference on Computer Supporter Cooperative Work*. pp. 175-186.
- 5 28. C. J. Van Rijsbergen. 1979. *Information Retrieval-Second Edition*. Butterworth & Co (Publisher) LTD.
29. G. Salton. and M.J. McGill. 1983. *Introduction to Modern Information Retrieval*. McGraw-Hill.
30. U. Shardanand and P. Maes. 1995. Social information filtering: Algorithms for automating "Word of Mouth". *ACM CHI'95 MOSAIC OF CREATIVITY*, pages 210-217.
- 10 31. B. D. Sheth. 1994. A Learning Approach to Personalized Information Filtering. Master's thesis, MIT Media Lab.
32. Study Group on Broadcasting System in the Multimedia Age (SGBSMA). 1995. Report by the Study Group on Broadcasting System in the Multimedia Age. Ministry of Posts and Telecommunication (MPT) of Japan, Press Release. [Http://www.mpt.go.jp/Report/Broadcasting-system/press-release.html](http://www.mpt.go.jp/Report/Broadcasting-system/press-release.html)
- 15 33. L. Terveen, W. Hill, B. Amento, D. MacDonald, and J. Creter. 1997. PHOAKS: A system for Sharing recommendations. *Communication of the ACM*, vol.40, no. 3, pages 59-62.
- 20 34. Universal Browser for Listening. 1997. <http://www.austinre.com/cgi/browser.cgi>.
35. J. Vassileva. 1994. A practical architecture for user modeling in a hypermedia-based information system. In *Proceedings of the 4th International Conference on user Modeling*, pages 115-120. MITRE Corporation.
- 25 36. J. M. Vidal and E. Durfee. 1995. Task planing agent in the UMDL. In *Proceedings of the 1995 Intelligent Information Agents Workshop*.
37. Video On Line. 1995. <http://www.vol.it/VOLB/browser.html>.
38. W. Wahlster. 1994. Toward personalized multimedia. In *3Proceedings of the 4th International Conference on User Modeling*, pages 3-4. MITRE Corporation.
- 30

39. D. Walley. 1996a. Data broadcasting vs. conventional on-line services. [Http://www.mbnet.mb.ca/~walleydc/advanta.html](http://www.mbnet.mb.ca/~walleydc/advanta.html).

40. D. Walley. 1996b. Free commercial mass-audience datacasting of popular information and software direct to personal computers.

5 41. [Http://www.mbnet.mb.ca/~walleydc/int_intr.html](http://www.mbnet.mb.ca/~walleydc/int_intr.html).

42. A. White and R. Jain. 1996. Algorithm and strategies for similarity retrieval. University of California, San Diego.

43. Bruce Krulwich; 1997; Lifestyle Finder, Intelligent user profiling using large scale demographic data.

10 44. G. Paliouras, C. Papatheodorou, V. Karkaletsis, C. Spyropoulos, and V. Malaveta; "Learning user communities for improving the services of information providers".

The disclosures of all references mentioned above and throughout the present specification are hereby incorporated herein by reference.

SUMMARY OF THE INVENTION

The present invention seeks to provide an improved scheduling and bandwidth allocation system for broadcasting, particularly but not exclusively suitable for data broadcasting.

5 The methods described above in the Herz et al patent do not appear to be suitable for a large system, such as is typical of most broadcasting systems. It would be impractical to provide an individual user profile and an individual virtual channel for each user. In effect, a unicast (one-to-one) system would be created in an implementation of the Herz et al system, rather than a broadcast
10 system, and the inherent economies of a broadcast system would thus be lost.

 The present invention addresses the problems of programming, scheduling, and bandwidth orientation, particularly but not only in a data broadcasting system. In the present invention, both user preferences and available bandwidth are taken into account. The present invention preserves the broadcast
15 nature of the system, rather than transforming the system into a unicast system, as would be the case if the prior art mentioned above were implemented.

 The present invention is particularly suited for broadcasting in an environment where more than one broadcast channel is used, and it is difficult or expensive to equip receivers with the ability to receive more than one channel
20 simultaneously. The term "receiver", as used throughout the present specification and claims, refers to any appropriate device adapted to receive an appropriate broadcast transmission. The problem addressed by the present invention is believed to be particularly acute in the case of data broadcasting, where broadcasting of items typically occurs in the background for use of the items
25 asynchronously, at a time typically later than the time of broadcast. The problem typically arises when a user of the system, equipped with a receiver, wishes to receive a plurality of items, comprising at least a first item and a second item which are broadcast on different channels, substantially simultaneously. In such a case, if the receiver utilized by the user is able to receive only a single channel at a
30 time, the user can not receive both the first item and the second item.

 The term "channel", as used throughout the specification and claims, is used in a broad sense to refer to any broadcasting path through which

transmission of information takes place. Without limiting the generality of the foregoing, the term "channel" may include, for example, any of the following: a standard allocated band of radio frequencies; a satellite transponder frequency; a defined portion of a standard allocated band of radio frequencies or a satellite
5 transponder frequency; and a single receivable stream of data within a multiplexed stream of data. It is appreciated, in particular, that the definition of "channel" includes a portion of the bandwidth of a standard allocated band or radio frequencies or a satellite transponder frequency, assigned to a specific use; and also includes a channel, as well-known in the art, allocated for broadcasting in any
10 appropriate wired system, including a cable system such as a cable television system, and any appropriate wireless system.

A goal of the present invention is to optimize the assignment of items to be broadcast to channels, such that each user will have the best possible chance of receiving the items in which the user is interested, even if the user's
15 receiver can only receive one channel at a time. Another goal of the present invention, preferably realized simultaneously with the previously mentioned goal, is to optimize the utilization of broadcast bandwidth. In the present invention, the aforementioned goals are accomplished by dividing users into communities of interest, herein also referred to as "communities", and matching items to
20 communities.

In a preferred embodiment of the present invention, user preferences are collected from a plurality of users. The users are divided into user communities based, at least in part, on user preferences and on available broadcast bandwidth, and a profile is defined for each community. Each of a plurality of
25 items to be broadcast is profiled according to the content of the item and scored according to how closely each item, as represented by the associated item profile, fits each community profile. One or more communities are assigned to each available broadcast channel. Items are then assigned for broadcast to a specific channel based, at least in part, on how closely the item profile fits the community
30 profile of each community sharing the channel.

There is thus provided in accordance with a preferred embodiment of the present invention a method for selecting an item to be broadcast in a

broadcasting system, the method including assigning a communication channel having a bandwidth to a community, for each one item of a plurality of available items, determining an affinity between the one item and the community and associating the affinity with the one item, computing a measure of remaining
5 bandwidth of the communication channel assigned to the community, and choosing at least one item to be broadcast from among the plurality of available items based, at least in part, on the affinity associated with the one item and on the measure of remaining bandwidth.

There is also provided in accordance with another preferred
10 embodiment of the present invention a method for selecting an item to be broadcast in a broadcasting system, the method including defining a plurality of communities, assigning a communication channel having a bandwidth to each one of the plurality of communities, for each one item of a plurality of available items determining an affinity between the one item and each one of the plurality of
15 communities, thereby producing a plurality of item-community affinities, and associating the plurality of item-community affinities with the one item, and for each one community of the plurality of communities computing a measure of remaining bandwidth of the communication channel assigned to the one community, and choosing at least one item to be broadcast from among the
20 plurality of available items based, at least in part, on the item-community affinity, associated with the one item, between the at least one item and the one community and on the measure of remaining bandwidth.

Further in accordance with a preferred embodiment of the present invention the choosing step includes re-computing the measure of remaining
25 bandwidth of the communication channel assigned to the community based on a measure of bandwidth required to broadcast the at least one item to be broadcast.

Still further in accordance with a preferred embodiment of the present invention the choosing step includes choosing in accordance with an affinity threshold associated with the one community.

30 Additionally in accordance with a preferred embodiment of the present invention the choosing step includes performing the following steps at least once: choosing, from among the plurality of available items, one item having the

highest item-community affinity of all items included in the plurality of available items, and removing the one item from the plurality of available items.

Moreover in accordance with a preferred embodiment of the present invention the method also includes modifying the affinity threshold based, at least
5 in part, on the measure of remaining bandwidth.

Further in accordance with a preferred embodiment of the present invention the method also includes allocating an updated bandwidth to the communication channel assigned to the community based, at least in part, on the measure of remaining bandwidth.

10 Still further in accordance with a preferred embodiment of the present invention the method also includes reassigning a new channel to the community based, at least in part, on the affinity threshold.

Additionally in accordance with a preferred embodiment of the present invention the method also includes producing a list of programming items
15 to be broadcast based, at least in part, on a result of the for each one item step and, at least in part, on a result of the for each one community step.

Moreover in accordance with a preferred embodiment of the present invention the method also includes producing a programming schedule based, at least in part, on a result of the for each one item step and, at least in part, on a
20 result of the for each one community step.

Further in accordance with a preferred embodiment of the present invention the method also includes providing at least one broadcast frequency, wherein at least one of the communication channels is assigned to the at least one broadcast frequency.

25 Additionally in accordance with a preferred embodiment of the present invention the at least one broadcast frequency includes a plurality of broadcast frequencies, and more than one of the communication channels is assigned to at least one of the plurality of broadcast frequencies.

Further in accordance with a preferred embodiment of the present invention the at least one broadcast frequency includes a satellite transponder
30 frequency.

There is also provided in accordance with another preferred embodiment of the present invention a bandwidth allocation method including receiving user information characterizing preferences of a plurality of broadcast item users, processing the user information to produce community information, the
5 community information including, for each one of a plurality of communities, community profile information, assigning a bandwidth to each one of the plurality of communities, assigning an item profile to each of a plurality of items to be broadcast, and choosing at least some of the plurality of items to be broadcast based, at least in part, on a measure of affinity between each of the plurality of
10 items to be broadcast and the community profile information for each one of the plurality of communities, and, at least in part, on the bandwidth assigned to each one of the plurality of communities.

Further in accordance with a preferred embodiment of the present invention the method also includes providing feedback including user profile
15 information, and varying, based, at least in part, on the feedback, the community profile information for at least one of the plurality of communities.

Still further in accordance with a preferred embodiment of the present invention the method also includes providing feedback including bandwidth utilization information, and varying, based, at least in part, on the
20 feedback, at least one of the following: the bandwidth assigned to at least one of the plurality of communities, and the measure of affinity between each of the plurality of items to be broadcast and the community profile information for at least one of the plurality of communities.

There is also provided in accordance with another preferred
25 embodiment of the present invention a system for selecting an item to be broadcast in a broadcasting system, the system including channel assignment apparatus for assigning a communication channel having a bandwidth to a community, affinity apparatus operative, for each one item of a plurality of available items, to determine an affinity between the one item and the community and associating the
30 affinity with the one item, bandwidth computation apparatus for computing a measure of remaining bandwidth of the communication channel assigned to the community, and item selection apparatus for choosing at least one item to be

broadcast from among the plurality of available items based, at least in part, on the affinity associated with the one item and on the measure of remaining bandwidth.

There is also provided in accordance with another preferred embodiment of the present invention a system for selecting an item to be broadcast in a broadcasting system, the system including community definition apparatus for
5 defining a plurality of communities, channel assignment apparatus for assigning a communication channel having a bandwidth to each one of the plurality of communities, affinity apparatus operative, for each one item of a plurality of available items, to determine an affinity between the one item and each one of the
10 plurality of communities, thereby producing a plurality of item-community affinities, and associating the plurality of item-community affinities with the one item, and community bandwidth apparatus operative, for each one community of the plurality of communities, to compute a measure of remaining bandwidth of the communication channel assigned to the one community and to choose at least one
15 item to be broadcast from among the plurality of available items based, at least in part, on the item-community affinity, associated with the one item, between the at least one item and the one community and on the measure of remaining bandwidth.

There is also provided in accordance with another preferred embodiment of the present invention a bandwidth allocation system including user
20 information apparatus operative to receive user information characterizing preferences of a plurality of broadcast item users, community information processing apparatus for processing the user information to produce community information, the community information including, for each one of a plurality of communities, community profile information, bandwidth apparatus for assigning a
25 bandwidth to each one of the plurality of communities, item profile apparatus for assigning an item profile to each of a plurality of items to be broadcast, and item choosing apparatus for choosing at least some of the plurality of items to be broadcast based, at least in part, on a measure of affinity between each of the plurality of items to be broadcast and the community profile information for each
30 one of the plurality of communities, and, at least in part, on the bandwidth assigned to each one of the plurality of communities.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood and appreciated more fully from the following detailed description, taken in conjunction with the drawings in which:

5 Fig. 1 is a simplified pictorial illustration of a broadcasting system constructed and operative in accordance with a preferred embodiment of the present invention;

 Fig. 2 is a simplified block diagram illustration of a preferred implementation of a portion of the apparatus of Fig. 1;

10 Fig. 3 is a simplified flowchart illustration of a preferred method of operation of the apparatus of Fig. 2;

 Fig. 4 is a simplified flowchart illustration of a preferred method of operation of a portion of the apparatus of Fig. 2;

15 Fig. 5 is a simplified flowchart illustration of a preferred method of operation of another portion of the apparatus of Fig. 2

 Fig. 6 is a simplified pictorial illustration of an agent-based user feedback collection system useful in understanding the preferred embodiment of Fig. 1;

20 Fig. 7 is a simplified flowchart illustration of a preferred method of operation of the system of Fig. 6;

 Fig. 8 is a simplified pictorial illustration of a preferred embodiment of a user profile, useful in understanding the method of Fig. 7;

 Fig. 9 is a simplified pictorial representation of a multimedia document representation, useful in understanding the method of Fig. 7; and

25 Fig. 10 is a simplified flowchart illustration of a preferred method of operation of a portion of step 1200 of Fig. 7.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

Reference is now made to Fig. 1 which is a simplified pictorial illustration of a broadcasting system, generally designated 100, constructed and operative in accordance with a preferred embodiment of the present invention.

5 The broadcasting system 100 comprises a broadcasting site 105 and a plurality, typically a large multiplicity, of user sites 110. For purposes of simplicity of description, six user sites 110 are shown in Fig. 1, but it is appreciated that any appropriate number of user sites 110, typically ranging from a small number such as six to a very large number such as several million or more, may be used.

10 The broadcasting system 100 also typically comprises a communications medium 115, preferably comprising a broadcast communications medium. The communications medium 115 is shown in Fig. 1 as a satellite-based communications medium, but it is appreciated that any appropriate broadcast communications medium known in the art may alternatively be used. Without
15 limiting the generality of the foregoing, it is appreciated that the choice of an appropriate communications medium would depend on the number of user sites 110 and on the distance between each of the user sites 110 and the broadcasting site 105, and that a range of typical examples of appropriate communications media which could be used would include the following: a wired communications
20 medium, such as a LAN, a cable television system, or a telephone system; a terrestrial wireless communications medium, such as a conventional terrestrial television broadcasting system or a digital terrestrial television broadcasting system; or a satellite-based communications medium.

The broadcasting site 105 preferably comprises broadcasting
25 equipment 120, which may comprise any appropriate conventional broadcasting equipment, it being appreciated that the broadcasting equipment 120 should be chosen for compatibility with the chosen type of communications medium 115. The broadcasting equipment 120 is typically operative to broadcast one or more items, as described further below, via the communications medium to the user sites
30 110.

The broadcasting site 105 also preferably comprises an item source 125. The item source 125 typically is implemented in a suitable combination of

hardware and software such as, for example, in a suitably programmed computer system preferably having very high capacity storage apparatus. Alternatively, the item source 125 may comprise any appropriate streaming source such as, for example, a digital camera delivering a live MPEG stream or other appropriate live data stream, or any other appropriate source of broadcast items. It is appreciated that the item source 125 may comprise any appropriate combination of item sources.

In a case where the item source 125 is operative to store items, the item source 125 is typically operative to store a plurality of items and to provide one or more of the stored items, typically upon external request, to a profiler and scheduler 130, also typically comprised within the broadcasting site 105. The profiler and scheduler 130 is typically implemented in a suitable combination of hardware and software. In a case where the item source 125 comprises a streaming source, as described above, the construction and operation of the system is similar, except that items are provided in streaming form rather than being provided from storage.

The profiler and scheduler 130 is preferably operative to receive a plurality of items from the item source 125 and user feedback 135 from the user sites 110 and to produce a list of items to be broadcast, per channel. The list of items to be broadcast, per channel, is preferably provided to the broadcasting equipment 120. A preferred implementation of the profiler and scheduler 130 is described more fully below with reference to Fig. 2.

It is appreciated that certain functions described herein as being performed by the profiler and scheduler 130, and hence within the broadcasting site 105, may alternatively be performed in a distributed manner, the results of the various functions then being transmitted to the broadcasting site 105 and processed therein in accordance with the description herein of the functioning of the profiler and scheduler 130. Without limiting the generality of the foregoing, it is appreciated that a distributed architecture utilizing user agents might be used in a preferred embodiment of the present invention. A particularly detailed description of certain aspects of a distributed architecture utilizing user agents, which

architecture might be utilized in a preferred embodiment of the present invention, is described below with reference to Figs. 6 - 10.

The operation of the apparatus of Fig. 1 is now briefly described. A first plurality of items is broadcast on at least one channel, typically on a second plurality of channels, via a communications medium 115 to user sites 110. At each user site over a period of time, one or more users chooses to receive and, typically, pay for one or more of the first plurality of items. Receiving and paying for items may be accomplished by utilizing a reception device 137, such as, for example, a television or a personal computer, located within each user site 110. Typically, each item is sent encrypted; methods for encrypting items, broadcasting the items, and making decryption of the items available only to those who have paid for the items are generally known as conditional access methods, and such methods are well-known in the art. A typical example of conditional access methods generally applicable to broadcasting and well-known in the art is described in US Patents 5,282,249 and 5,481,609 to Cohen et al. An example of conditional access techniques particularly suited to a data broadcasting context is described in the Key Delivery application.

Individual user feedback information 140, comprising information on the items chosen and/or other preferences of each user, is collected by each reception device. The individual user feedback information 140 may be individually transmitted from each user site 110, using any appropriate method well-known in the art such as a telephone return path or a satellite return path, to the profiler and scheduler 130. Alternatively, as shown in Fig. 1, the individual user feedback information 140 may be aggregated using any appropriate aggregation apparatus (not shown) implementing any appropriate aggregation method, to produce user feedback 135, the user feedback 135 being transmitted to the profiler and scheduler 130.

A distributed architecture utilizing user agents, which architecture may be useful in a preferred embodiment of the aggregation apparatus (not shown), is described below with reference to Figs. 6 - 10. It is appreciated, however, that a wide variety of other embodiments, including other preferred

embodiments, may be possible, including embodiments which may or may not utilize user agents.

The profiler and scheduler 130 preferably uses the received user feedback information 135 and the items received from the item source 125 to produce an item list per channel. The item list per channel is preferably produced in such a way as to maximize the likelihood that, at any user site 110, the reception device 137 may receive items preferred according to the individual user feedback information from a single channel at any one time. In other words, the profiler and scheduler 130 strives to maximize the likelihood that even if tuning from channel to channel may be necessary to receive all preferred items, it will not generally be the case that it would be necessary to tune to more than one channel simultaneously in order to receive all preferred items.

Preferably, each reception device 137 is provided with time and channel information for receiving each preferred item, the time and channel information being provided to each reception device 137 using any means known in the art such as, for example, as part of an electronic program guide to available items.

Reference is now made to Fig. 2, which is a simplified block diagram illustration of a preferred implementation of the profiler and scheduler 130 of Fig. 1. Generally, the elements of the apparatus of Fig. 2 may be implemented in an appropriate combination of hardware and software, as is well known in the art.

It is appreciated, as described above, that the elements of the apparatus of Fig. 2, as described below, may alternatively be distributed in portions of the system of Fig. 1 other than the profiler and scheduler 130. For the sake of simplicity of description and without limiting the generality of the foregoing, the elements of Fig. 2 will be described herein as if they were comprised in the profiler and scheduler 130.

The apparatus of Fig. 2 comprises a user profiler 145. The user profiler 145 typically receives user information, typically user information comprised in the user feedback 135 of Fig. 1, from a plurality of users. The user

profiler 145 is preferably operative to produce user profiles, typically comprising aggregated user profiles describing groups of users with similar interests.

The apparatus of Fig. 2 also comprises a clusterer 150. The clusterer 150 typically receives the user profiles produced by the user profiler 145 and also receives channel information, either from another component of the system of Fig. 1 (not shown) or from a human operator via any appropriate input device or method. The channel information typically comprises a number of channels allocated for broadcasting and, for each channel, the bandwidth assigned to that channel. Alternatively, each channel may have the same bandwidth as every other channel, in which case only one bandwidth need be supplied in the channel information, but allowing the possibility of different bandwidths for different channels is believed to be preferred.

The clusterer 150 is preferably operative, based on the input received as described above, to produce community information, the community information comprising a number of communities and a profile for each community. Typically, each profile comprises a set of characteristics of items likely to be preferred by the community associated with the profile. For example, the profile may comprise a list of keywords, which keywords may appropriately be associated with one or more items. It is appreciated that other types of characteristics may also be used.

Preferably, in producing community information, the cluster 150 is operative to employ an appropriate optimization algorithm to optimize the degree of commonality of interests between the members of each community, such that, preferably, the members of each community will tend to prefer to receive a given set of items having given characteristics. An example of a preferred method for producing such communities is described in detail below with reference to Figs. 6 - 10, but it is appreciated that other methods may also be used.

The clusterer 150 is also operative to produce a bandwidth assignment per community and a community / channel mapping.

The following is one particular example of possible values of the possible inputs and outputs of the cluster 150; the example is offered for

clarification of the above description only, is not meant to be limiting, and is chosen to be compatible with the simplified example of Fig. 1:

Inputs

User profiles

Number of channels: 3

Bandwidth for channel 1: 2 Mbit/sec

Bandwidth for channel 2: 1 Mbit/sec

Bandwidth for channel 3: 1 Mbit/sec

Outputs

Number of communities: 5

Community 1: Sports (with details of sports profile)

Community 1: 1.5 Mbit/sec

Community 1: channel 1

Community 2: News (with details of news profile)

Community 2: 0.5 Mbit/sec

Community 2: channel 1

Community 3: Business (with details of business profile)

Community 3: 1 Mbit/sec

Community 3: channel 2

Community 4: Movies (with details of movies profile)

Community 4: 0.5 Mbit/sec

Community 4: channel 3

Community 5: Fashion (with details of fashion profile)

Community 5: 0.5 Mbit/sec

Community 5: channel 3

5 In the above example, the Sports and News communities have been assigned to channel 1; the Business community to channel 2; and the Movies and Fashion communities to channel 3. The names of communities, such as “Sports” and “News” are used herein in order to simplify the description of the present invention, it being appreciated that a community may comprise users having profiles of varying complexity which may not be capable of description by a

simple term such as "Sports" or "News", the community typically representing a mixture of areas of interest having appropriately similar weights designating interests. In the present example, however, it will be seen that a user who is primarily interested in News type items and secondarily interested in Sports type items may remain tuned to channel 1 and receive most of the items in which they are interested. It is also appreciated that, using apparatus and methods well known in the art, preferred channel tuning information may be communicated to each user, such that each user need not know to which channel they should be tuned.

It is appreciated that the example provided above is simplified for ease of description, and that much more complicated examples would typically be used in an actual system. In particular, and without limiting the generality of the foregoing, it is appreciated that there may be a time dimension in the community information. For example, the Sports community typically prefers to receive broadcast items in the evening hours, while the Business community typically prefers to receive broadcast items in the morning hours. As a further example, the Sports community might prefer to receive broadcast items on weekends, while the Business community might prefer to receive broadcast items on weekdays. It is appreciated, therefore, that any further use of community information, as described herein, would preferably take into account the time of day and/or day of the broadcast in deciding which items to schedule for broadcast.

The apparatus of Fig. 2 also comprises an item profiler 155. The item profiler 155 typically receives as input a plurality of items and assigns to each item an item profile, the plurality of item profiles being produced as output by the item profiler.

In a first preferred embodiment of the present invention, each item received as input by the item profiler 155 includes a list of characteristics, such as keywords, characterizing the item. In the first preferred embodiment at least some characteristics in the list of characteristics may be associated with a particular portion of the item. For example, if an item comprises a news program, the first 5 minutes might be associated with the keyword "international", the second 5 minutes with the keyword "business", and the last 10 minutes with the keyword "sports". In the first preferred embodiment, the profiler may be operative to

produce the item profile by assigning weights to each input keyword. In the previous example, the profile might be assigned as follows:

	Keyword	Weight
	international	25%
5	business	25%
	sports	50%

It is appreciated that more complex keyword and profile assignments are possible, with a single portion of the item being assigned a plurality of keywords, and that the example given is not meant to be limiting.

10 In a second preferred embodiment of the present invention, each item received as input by the item profiler 155 is automatically analyzed and an item profile is generated for each item based, at least in part, on a result of analyzing the item. Systems and methods for automatically analyzing classifying a document are well-known in the art, and include textual parsers and agent-based
15 classification systems such as, for example: Autonomy, commercially available from Autonomy Systems Limited, Cambridge, UK; and Firefly, commercially available from Agents Inc., Cambridge, Massachusetts, USA.

Typically, each item profile also comprises other item characterizing information, preferably including a measure of size, such as size in megabytes, of
20 the associated item.

The apparatus of Fig. 2 also preferably includes an assigner/thresholder 160. The assigner/thresholder 160 preferably receives as input: the output of the clusterer 150, comprising the number of communities and a community profile of each community, as described above; and the output of the
25 item profiler 155, comprising a plurality of item profiles, as described above. The assigner/thresholder 160 is operative to determine, for each of the plurality of items whose item profile is received as input, the affinity of the item to each of the plurality of communities. The affinity between an item and a community may be determined using any appropriate measure of affinity or closeness, typically a
30 measure of affinity between components, such as keywords, of the item profile and components, such as keywords, of each community profile. One example of an appropriate measure of affinity is described below with reference to Figs. 6 - 10,

but it is appreciated that other methods, such as those mentioned above in regard to automatically classifying a document, may also be used.

Typically, the assigner/thresholder 160 is operative to use the affinity of each item to assign each item to one or more communities. The
5 assigner/thresholder may use any appropriate method for assigning each item to one or more communities. In one possible example of a typical method for assigning each item to one or more communities, the assigner/thresholder 160 may assign a threshold affinity value for each community, and may assign to each community all items whose affinity to that community exceed the threshold value,
10 it being appreciated that any one item may thus be assigned to more than one community.

Preferably, the assigner/thresholder 160 produces as output a list, per community, of items to be scheduled for that community, based on the assigning of items to each community as described above, the list typically
15 comprising an indication of the level of affinity between each item and that community.

A preferred method of operation of the assigner/thresholder 160 is described in more detail below with reference to Fig. 4.

The apparatus of Fig. 2 also preferably comprises a scheduler 165.
20 The scheduler 165 typically receives as input the list of items to be scheduled per community, output as described above by the assigner/thresholder 160. The scheduler 165 may also receive as input a preferred item schedule from an external source (not shown), such as, for example, from user input of a human being monitoring the operation of the apparatus of Fig. 2, via any appropriate user input
25 means and/or method. The preferred item schedule, if received, may typically indicate that a particular item is to be broadcast no more than a first specified number of times each day, and/or no fewer than a second specified number of times each day.

The scheduler 165 also typically receives as input the bandwidth per
30 channel and the community / channel mapping produced as output by the clusterer 150, as described above.

The scheduler 165 is typically operative to process the received input, as described above, and to produce as output an item list per channel, which is typically provided to broadcasting equipment such as the broadcasting equipment 120 of Fig. 1. The scheduler 165 may use any appropriate method for producing the item list per channel.

In one possible example of a typical method of producing an item list per channel, for each community assigned to each channel the scheduler 165 may keep track of a remaining bandwidth for the next predetermined period of time. The scheduler 165 may then assign to that channel the next item for that community as found in the list of items received as input, and may decrement the remaining bandwidth for that channel in accordance with the size of the item to be broadcast. The scheduler may thus continue until the assigned bandwidth for the one community for the next predetermined period of time is exhausted, and then proceed to the next community until all bandwidths of all communities are exhausted. The scheduler 165, in assigning items, may assign the item with the greatest affinity for the community first, and then proceed in order with items of lesser affinity.

The scheduler 165 may also produce as output a feedback for scheduler adjustment directed back to the assigner/thresholder 160. Typically, such a feedback would be produced as follows: if not all bandwidth for a given community is used in a given period of time, the feedback would indicate that the affinity threshold for that community should be lowered, thus tending to produce more items for that community in the future; and, if an excessive number of items for a given community is received, such that not all items can be scheduled in a reasonable period of time or such that it is difficult to fulfill a preferred item schedule, the feedback would indicate that the affinity threshold for that community should be raised, thus tending to produce fewer items for that community in the future.

The scheduler 165 may also produce as output a feedback for channel/bandwidth adjustment directed back to an apparatus (not shown) for channel/bandwidth adjustment or to a human operator, a decision based on said feedback being typically directed to the clusterer 150. Typically, such a feedback

would be produced as follows: if not all bandwidth for a given community is used in a given period of time, the feedback would indicate that the bandwidth for that community should be lowered; and, if an excessive number of items for a given community is received, such that not all items can be scheduled in a reasonable period of time or such that it is difficult to fulfill a preferred item schedule, the feedback would indicate that the bandwidth for that community should be raised. Optionally, the feedback may also include an indication that channel/bandwidth assignments should be changed in order to accommodate the changes in bandwidth assignments to communities by the clusterer 150.

10 A preferred method of operation of the scheduler 165 is described in more detail below with reference to Fig. 5.

Reference is now made to Fig. 3, which is a simplified flowchart illustration of a preferred method of operation of the apparatus of Fig. 2. The method of Fig. 3 preferably includes the following steps:

15 A plurality of communities is defined (step 170). Each community is assigned a communications channel having a defined bandwidth (step 175).

For each one item of a plurality of items available for broadcast: an affinity is determined between the one item and each one of the plurality of communities, thus producing a plurality of item-community affinities (step 180); and the plurality of item-community affinities is associated with the one item (step 185).

For each one community of the plurality of communities: a measure of remaining bandwidth of the communication channel assigned to the one community is computed (step 190); and at least one item to be broadcast is chosen, from among the plurality of available items based, at least in part, on the item-community affinity, associated with the one item, between the at least one item and the one community, and on the measure of remaining bandwidth (step 195).

Reference is now made to Fig. 4, which is a simplified flowchart illustration of a preferred method of operation of the assigner/thresholder 160 of Fig. 2. The method of Fig. 4 preferably comprises the following steps, which are typically performed once for each item to be assigned, except as otherwise described below:

Input is typically received as follows:

- from the item profiler, an item profile for one item. Alternatively, an item profile for a plurality of items may be received, in which case only one item profile at a time is preferably processed by the method of Fig. 4; and

5 - from the clusterer, community profiles for a plurality of communities, shown for simplicity of description in Fig. 4 as communities 1 through n.

The following steps are preferably performed for each community i:

10 The item profile is matched against community profile i (step 200), and a score i is generated indicating the degree of match between the item profile and community profile i.

 A check is made as to whether, for a current time interval j, the score i is greater than the threshold for community i during time interval j, symbolized by threshold ij (step 210). If not the next community is considered
15 (step 220). If so, the item is assigned to community i (step 230), and then the next community is considered (step 220).

 After all communities have been considered, an output list of communities to which the item is assigned is generated and sent to the scheduler (step 240).

20 At any appropriate time, the scheduler may generate feedback, as described above with reference to Fig. 2; the feedback is received and a new threshold for community i in time interval j is computed accordingly (step 250), with the result being supplied to step 210, described above. It is appreciated that any appropriate means may be used by step 250 to compute a new threshold, and
25 that many types of feedback are possible, for which many possible types of threshold computation would be appropriate, respectively.

 Reference is now made to Fig. 5, which is a simplified flowchart illustration of a preferred method of operation of the scheduler 165 of Fig. 2.

 The method of Fig. 5 typically receives the following items as input:

30 - items to be scheduled per community, including a match of each item to a community i at time interval j;
 - a preferred item schedule; and

- a list of bandwidth assignments per community at time interval j.

The method of Fig. 5 preferably includes the following steps, which are typically performed once for each item received:

5 A determination is made as to whether there is sufficient vacant bandwidth in community i at time interval j to fit the requested item into the schedule (step 260). If the bandwidth is sufficient (step 270), processing continues with step 280, described below; otherwise, processing continues with step 310, described below.

10 A check is made as to whether there is abundant extra bandwidth (step 280). If there is abundant extra bandwidth, the threshold for community i at interval j is lowered (step 290), thus allowing more items to be assigned to community i in the future, during interval j.

The item is assigned to a broadcast channel k associated with community i (step 300).

15 Continuing from step 270, the threshold for community i during interval j is raised (step 310), thus allowing fewer items to be assigned to community i in the future, during interval j.

A check is made as to whether the item matches any other community other than community i (step 320); if so, processing continues, with a
20 new value assigned to i to represent the next such community, at step 260.

If not, a check is made as to whether the community threshold was raised or lowered past a predefined limit (step 330). The predefined limit generally represents a limit, typically set by a human being, indicating the highest threshold, indicating that many potential items are being excluded, and/or the
25 lowest threshold, indicating that many marginally related items are being included, which is acceptable during operation of the system. Typically, there will be both a highest, or upper, threshold and a lowest, or lower threshold. If a community threshold was neither raised nor lowered past the predefined threshold, processing continues with step 350, described below.

30 If a community threshold was raised or lowered past the predefined threshold, feedback is output indicating that bandwidth and possibly channel assignment for community i should be adjusted (step 340). Typically, the

adjustment of step 340 comprises bandwidth being adjusted upward if the threshold was raised past the predefined upper threshold, and with bandwidth being adjusted downward if the threshold was lowered past the predefined lower threshold. A channel assignment might be adjusted based on unavailability of
5 enough bandwidth on a given channel for the bandwidth to be adjusted upward, or based on efficiency of channel utilization if the community is transferred to another channel, in case of the bandwidth being adjusted downward; or on any other appropriate consideration of combination of considerations.

The item is then scheduled for broadcast on all assigned channels
10 (step 350), with the assignments typically being sent to broadcast equipment.

Reference is now made to Fig. 6 which is a simplified pictorial illustration of an agent-based user feedback collection system useful for understanding the preferred embodiment of Fig. 1. The system of Fig. 6 comprises a plurality of user stations 1100, each of the plurality of user stations 1100 being
15 operatively associated with a user profile agent 1110. It is appreciated that, typically, a multiplicity of user stations 1100, such as, for example, 1 million or more user stations 1100, is provided, a plurality of user stations 1100 being shown in Fig. 1 for sake of simplicity of description. It is also appreciated, as also described below, that other agents (not shown) in addition to the user profile agent
20 1110, may also be operatively associated with each of the plurality of user stations 1100.

Each of the user stations 1100 preferably comprises an appropriate computing system, such as a multimedia computing system as is well-known in the art, equipped and operative to receive broadcasts, typically comprising multimedia
25 data broadcasts, from a broadcast center 1120. It is appreciated, as is well known in the art, that the broadcast center 1120 may broadcast using any one or more of a variety of appropriate broadcast methods, including conventional terrestrial broadcast, digital terrestrial broadcast, satellite broadcast, cable broadcast, or any other appropriate broadcast method. It is further appreciated that each user station
30 1100 is preferably equipped to receive broadcasts of a type broadcast by the broadcast center 1120.

Each of the user profile agents 1110 is preferably implemented in a combination of computing hardware and software, as is well known in the art, and is preferably operative to observe the behavior of a user of the associated user station 1100 and to create a user profile, as described more fully below, characterizing the preferences of the user of the associated user station 1100.

The system of Fig. 6 also comprises a user profile subsystem 1130. The user profile subsystem 1130 is preferably implemented in an appropriate combination of computing hardware and software, as is well known in the art, and is typically located remotely from each of the user stations 1100, and in operative communication therewith. The user profile subsystem 1130 is operative to receive a user profile from each of the user profile agents 1110, to analyze the received profiles, and to derive therefrom a plurality of community profiles characterizing the multiplicity of user stations 1100 in accordance with a plurality of communities, each community preferably sharing a set of preferences and characteristics. Typically, the user profile subsystem 1130 produces an output comprising community information, preferably in the form of a plurality of community profiles each characterizing one community.

The user profile subsystem 1130 typically comprises a plurality of agents, as described below. The plurality of agents comprised within the user profile subsystem 1130 may, when referred to generally, be referred to as "integrated user profile agents". For simplicity of description, the integrated user profile agents are collectively referred to herein as the "user profile subsystem 1130", it being appreciated that, because of possible separate implementation and possible geographic dispersal of the integrated user profile agents, the term "user profile subsystem 1130" is itself a term which has been adopted mainly for the sake of simplicity of description.

The user profile subsystem 1130 typically comprises a plurality of clustering agents 1140, each clustering agent 1140 being operatively associated with a subset of the user stations 1100 and preferably being operative to produce one or more preliminary community profiles based on information received from the associated user profile agents 1110 associated with each of the user stations 1100 in the subset of the user stations 1100.

The user profile subsystem 1130 also typically comprises at least one union set agent 1150. The union set agent 1150 is preferably operative to cooperate with a plurality of clustering agents 1140 and to adjust the operation thereof in such a way that community information produced by each clustering agent 140 is compatible with, and tends to form the same communities as, that
5 agent 140 is compatible with, and tends to form the same communities as, that produced by each other clustering agent 1140. Specifically, and without limiting the generality of the foregoing, the union set agent 1150 is preferably operative to examine community profiles produced by a plurality of clustering agents 1140, to identify similar community profiles produced by different clustering agents 1140,
10 to unify the similar community profiles, and to report the unified profiles back to each of the different clustering agents 1140 so that the different clustering agents 1140 will operate with the same community profiles.

The system of Fig. 6 also comprises at least one community agent 1155, typically comprising a plurality of community agents 1155. Each
15 community agent 1155 is preferably associated with one community profile produced by one or more clustering agents 1140 as described above, and is preferably operative to aggregate information relating to the one community profile as reported by the one or more clustering agents 1140.

The system of Fig. 6 also comprises a broadcasting agent 1160,
20 which is typically implemented in an appropriate combination of computer hardware and software. The broadcasting agent 1160 preferably receives the community information produced by the user profile subsystem 1130 and is preferably operative, based on the received community information, to modify the operation of the broadcast center 1120 to take into account community preference
25 information comprised in the community information. By way of example only and without limiting the generality of the foregoing, modifying the operation of the broadcast center 1120 may include modifying the content of future broadcasts.

The operation of the apparatus of Fig. 6 is now briefly described. The broadcast center 1120 broadcasts a plurality of broadcast items to the
30 multiplicity of user units 1100. The user of each user unit 1100 elects to receive one or more broadcast items; it is appreciated that, in some cases, a user may elect not to receive any items. It is appreciated, as is well known in the art, that a user

may elect to receive a broadcast item using any appropriate technique, such as, for example: the user may select an individual item or a group of items in advance; the user may select an individual item or a group of items on demand; the user may program the user unit 1100 to select a particular item, an item matching certain
5 criteria, or an item broadcast at a particular time; or any appropriate combination of the above, including a combination with other techniques well known in the art. It is also appreciated that a filter agent (not shown) may automatically download, accept, or record only those items which are deemed to be of interest to the user based on criteria determined by the filter agent, by other agents, or otherwise,
10 particularly as described below.

Each user profile agent 1110 is preferably operative, as described more fully below, to create a user profile describing the item-selection preferences of the user of the associated user unit 1100. Each user profile agent 1110 then passes on the user profile to the user profile subsystem 1130, specifically to one of
15 the clustering agents 1140. Each clustering agent 1140 is preferably operative, as described more fully below, to cluster information from a plurality of user profiles and to create therefrom community profile information, representing a clustered aggregate of information about preferences of a plurality of users.

Preferably, the union set agent 1150 is operative, as described more
20 fully below, to coordinate the activities of the clustering agents 1140 in order to assure that uniform community profile information, representing clustered information across all of the user units 1100, is produced. Community information from a plurality of clustering agents 1140 is then typically aggregated into a community profile for each community by one or more community agents 1155,
25 each community agent 1155 preferably being associated with one community profile. The community information, typically in the form of the community profiles, is then passed on the broadcasting agent 1160, where it is used as input in broadcast decisions of the broadcast center 1120, including decisions as to the content of future broadcasts.

30 Reference is now made to Fig. 7, which is a simplified flowchart illustration of a preferred method of operation of the system of Fig. 6. The method of Fig. 7 preferably comprises the following steps:

A user profile is derived for each user of a multiplicity of data item users (step 1170). It is appreciated that a wide variety of different methods, including methods known in the art, may be used to derive a user profile, and that a user profile may have any one of a wide variety of structures, including structures which are well known in the art. It is further appreciated that not every user need have exactly the same user profile structure as every other user.

Reference is now additionally made to Fig. 8, which is a simplified pictorial illustration of a preferred embodiment of a user profile, useful in understanding the method of Fig. 7. In the example user profile of Fig. 8, a profile is depicted as a tree, in which each level represents a different level of information about user preferences, each node in the tree being represented as a vector, the sum of whose component values is 1.0. Reference is now additionally made to Fig. 9, which is a simplified pictorial representation of a multimedia document representation, useful in understanding the method of Fig. 7. It is appreciated that the example document representation of Fig. 9 is generally compatible with the example user profile of Fig. 8.

Referring back to Fig. 7, user profile information from at least some of the multiplicity of data item users is clustered together to produce preliminary community information (step 1180). The clustering in step 1180 may use any appropriate clustering method for combining together similar profiles. Preferably, an appropriate clustering method may be chosen from the field of information retrieval technology, with the additional feature that more than one level of user profile, such as more than one level of the example user profile of Fig. 8, is used for clustering.

Clustering of user profiles preferably is based on similarity between user profiles. Similarity between two user profiles, based on user classes and subclasses, may preferably be computed in accordance with the following formula:

$$\text{similarity}(\text{user1}, \text{user2}) = \frac{\sum_{C} \sum_{Si} \text{user1_classi} \cdot \text{user1_subclassij} \times \text{user2_classi} \cdot \text{user2_subclassij}}{\sqrt{\sum_{C} \sum_{Si} (\text{user1_classi} \cdot \text{user1_subclassij})^2} \cdot \sqrt{\sum_{C} \sum_{Si} (\text{user2_classi} \cdot \text{user2_subclassij})^2}}$$

where user1 and user2 are class vectors of two user profiles, each user profile having a plurality of subclasses, C being the number of classes and Si the corresponding number of subclasses. It is appreciated that the above formula comprises one example of a suitable formula, and that other formulas may be used.

5 Preferably, clustering is carried out using an appropriate agglomerative clustering method, as is well known in the art. One example of an appropriate method for clustering is as follows:

0. Identify each user profile as a point in the space to be merged.
1. Initialize each point as a cluster.
- 10 2. Find the two closest clusters in the space and merge them together.
3. Compute the distance from the newly formed cluster to all other clusters.
4. If more than one cluster remains, return to step 2.

15 Typically, a method such as that described above outputs a complete binary tree wherein each internal node of the tree represents a merge between two cluster and the leaves are the points in the space, or user profiles. Typically, the lower levels in the tree are small tight clusters while the higher levels are larger, general clusters. The indicated method, which operates on the available data and
20 is not limited by external parameters, is considered, based on the work of the inventors of the present invention, to be preferred in the present invention.

In the method described above, a measure of closeness or distance is important in determining which clusters to merge. A number of ways of determining distance are well known in the art. In the present invention, it is
25 believed to be preferred to use a "complete link" method, in which the distance between two clusters is the distance between their two furthest points. This method is believed likely to produce bound clusters, as opposed to long, "straggly" clusters.

An approach to expressing the method described above is to use the
30 Lance Williams dissimilarity update formula, which is well-known in the art. The formula enables incremental recalculation of cluster distances. The Lance-Williams dissimilarity update formula is given as follows:

$d(c_i, j, c_k) = \alpha_i \cdot d(c_i, c_k) + \alpha_j \cdot d(c_j, c_k) + \beta \cdot d(c_i, c_j) + \gamma |d(c_i, c_k) - d(c_j, c_k)|$
 where, for the preferred case of complete link, the following parameters are preferably used:

$$\alpha_i = \frac{1}{2}$$

$$\beta = 0$$

$$\gamma = \frac{1}{2}$$

The preliminary community information is modified to produce community information (step 1190). Referring back to Fig. 6, it is appreciated that each clustering agent 1140 is operative to produce preliminary community information relating to a plurality of associated user stations 1100, which plurality of associated user stations 1100 is generally disjoint from each other plurality of user stations 1100 associated with another clustering agent 1140. For reasons of efficiency, it is desirable nonetheless for the preliminary community information produced by each clustering agent 1140 to be modified and made uniform with community information produced by other clustering agents 1140. Preferably, the union set agent 1150 is primarily operative to carry out step 1190, with final agglomeration at the community level being carried out by the community agent 1155, as described above.

Referring back to Fig. 7, in step 1190, preferably preliminary community information, preferably in the form of preliminary community profiles produced by a plurality of clustering agents, is unified preferably using methods similar to those described above for user profiles in step 1180. Preferably, as described above with reference to Fig. 6, the union set agent sends information back to the clustering agents, the information comprising updated community profile information which replaces the preliminary community profile information. In this way, as described above with reference to Fig. 6, different clustering agents will produce information agglomerated into the same community profiles. In this way, since each community agent is associated with one community profile, the union set agents preferably determine the number of community profiles and the number of community agents.

The community information is then utilized for modifying a broadcast schedule (step 1200). Typically, the broadcasting center has a list of items to broadcast, the list being made available to a broadcasting agent. It is appreciated that each item, sometimes referred to herein as a “document”, may
 5 comprise any article or item typically broadcast by a broadcaster, such as, for example, a document, an image, a video, a computer program, a television program, an update to any of the previously listed items, or any other appropriate article or item. The agent attempts to prepare, based on document information and on community profiles, a broadcasting program which will maximize the
 10 satisfaction of the user community. Preferably, this task comprises extracting a multimedia document representation based on information that accompanies the multimedia document; and scoring all documents and preparing a broadcast program for each community based on the results of the scoring.

The broadcasting agent preferably determines the similarity between
 15 the community profiles and the document collection. Different communities rate the documents differently. Therefore, for each document in the document collection, the broadcasting agent prepares a vector of scores with respect to each community profile. In other words, for each document we have a vector of scores where score s_i is the score that the document received by community i . A
 20 preferable function that the broadcasting agent may use for scoring the documents is as follows:

$$\begin{aligned} & \text{Score}_{\text{community_profile, document}} = \\ & \text{Community_Class_Weight}_c \times \\ & \text{Community_Subclass_Weights} \times \\ & \left[\sum_D \text{Com_prof_desc}_d \times \left\{ \frac{\sum_k \text{doc_term}_i \times \text{prof_term}_i}{\sqrt{\sum_k \text{doc_term}_i^2 \times \sum_k \text{prof_term}_i^2}} \right\} \right] \end{aligned}$$

where:

Community_Class_Weight_c is the weight of class c according to the community profile;

Community_Subclass_Weights is the weight of subclass s of class c of the community profile;

5 Com_prof_descd is the description field's weight of the community profile;

D is the number of description fields in the description vector of the profile; and

K is the number of terms in the term vector.

10 A broadcast program is then chosen for each community. Reference is now additionally made to Fig. 10, which is a simplified flowchart illustration of a preferred method of operation of a portion of step 200 of Fig. 2. The method of Fig. 10 comprises a preferred method for choosing a broadcast program for each community. The method of Fig. 10 is self explanatory, except as described below.

15 Preferably, after an initial community program is prepared for each community, an attempt is made to extend the programs to be broadcast by scanning the list of documents selected to be broadcast for all users and assigning, to each community, additional relevant documents that were selected to be broadcast but were not assigned to that community program. In this context, a relevant
20 document for a community is preferably defined as a document whose score for that community exceeds a relevancy threshold, the relevancy threshold typically comprising a parameter pre-defined by an operator of the system.

It is appreciated that various features of the invention which are, for clarity, described in the contexts of separate embodiments may also be provided in
25 combination in a single embodiment. Conversely, various features of the invention which are, for brevity, described in the context of a single embodiment may also be provided separately or in any suitable subcombination.

It will be appreciated by persons skilled in the art that the present invention is not limited by what has been particularly shown and described
30 hereinabove. Rather the scope of the invention is defined only by the claims which follow:

What is claimed is:

CLAIMS

1. A method for selecting an item to be broadcast in a broadcasting
5 system, the method comprising:
 assigning a communication channel having a bandwidth to a
community;
 for each one item of a plurality of available items, determining an
affinity between said one item and the community and associating said affinity
10 with said one item;
 computing a measure of remaining bandwidth of the communication
channel assigned to the community; and
 choosing at least one item to be broadcast from among said plurality
of available items based, at least in part, on said affinity associated with said one
15 item and on said measure of remaining bandwidth.
2. A method for selecting an item to be broadcast in a broadcasting
system, the method comprising:
 defining a plurality of communities;
20 assigning a communication channel having a bandwidth to each one
of said plurality of communities;
 for each one item of a plurality of available items:
 determining an affinity between said one item and each one
of said plurality of communities, thereby producing a plurality of item-community
25 affinities, and associating said plurality of item-community affinities with said one
item; and
 for each one community of said plurality of communities:
 computing a measure of remaining bandwidth of the
communication channel assigned to the one community; and
30 choosing at least one item to be broadcast from among said
plurality of available items based, at least in part, on said item-community affinity,

associated with said one item, between said at least one item and said one community and on said measure of remaining bandwidth.

3. A method according to claim 2 and wherein said choosing step
5 comprises re-computing the measure of remaining bandwidth of the communication channel assigned to the community based on a measure of bandwidth required to broadcast said at least one item to be broadcast.

4. A method according to claim 2 or claim 3 and wherein said
10 choosing step comprises choosing in accordance with an affinity threshold associated with said one community.

5. A method according to claim 4 and wherein said choosing step
comprises performing the following steps at least once:

15 choosing, from among said plurality of available items, one item having the highest item-community affinity of all items comprised in said plurality of available items; and

removing said one item from said plurality of available items.

20 6. A method according to claim 4 or claim 5 and also comprising modifying said affinity threshold based, at least in part, on said measure of remaining bandwidth.

7. A method according to any of claims 2 - 6 and also comprising
25 allocating an updated bandwidth to said communication channel assigned to said community based, at least in part, on said measure of remaining bandwidth.

8. A method according to any of claims 4 - 7 and also comprising
30 reassigning a new channel to said community based, at least in part, on said affinity threshold.

9. A method according to any of claims 2 - 8 and also comprising:

producing a list of programming items to be broadcast based, at least in part, on a result of the for each one item step and, at least in part, on a result of the for each one community step.

5 10. A method according to any of claims 2 - 9 and also comprising:
producing a programming schedule based, at least in part, on a result of the for each one item step and, at least in part, on a result of the for each one community step.

10 11. A method according to any of claims 2 - 10 and also comprising:
providing at least one broadcast frequency, wherein at least one of said communication channels is assigned to said at least one broadcast frequency.

12. A method according to claim 11 and wherein said at least one
15 broadcast frequency comprises a plurality of broadcast frequencies, and
more than one of said communication channels is assigned to at least one of said plurality of broadcast frequencies.

13. A method according to claim 11 and wherein said at least one
20 broadcast frequency comprises a satellite transponder frequency.

14. A bandwidth allocation method comprising:
receiving user information characterizing preferences of a plurality of broadcast item users;
25 processing the user information to produce community information, the community information comprising, for each one of a plurality of communities, community profile information;
assigning a bandwidth to each one of the plurality of communities;
assigning an item profile to each of a plurality of items to be
30 broadcast; and
choosing at least some of the plurality of items to be broadcast based, at least in part, on a measure of affinity between each of the plurality of

items to be broadcast and the community profile information for each one of the plurality of communities, and, at least in part, on the bandwidth assigned to each one of the plurality of communities.

5 15. A method according to claim 14 and also comprising:
 providing feedback comprising user profile information; and
 varying, based, at least in part, on the feedback, the community
profile information for at least one of the plurality of communities.

10 16. A method according to claim 14 or claim 15 and also comprising:
 providing feedback comprising bandwidth utilization information;
and
 varying, based, at least in part, on the feedback, at least one of the
following: the bandwidth assigned to at least one of the plurality of communities;
15 and the measure of affinity between each of the plurality of items to be broadcast
and the community profile information for at least one of the plurality of
communities.

17. A system for selecting an item to be broadcast in a broadcasting
20 system, the system for selecting comprising:
 channel assignment apparatus for assigning a communication
channel having a bandwidth to a community;
 affinity apparatus operative, for each one item of a plurality of
available items, to determine an affinity between said one item and the community
25 and associating said affinity with said one item;
 bandwidth computation apparatus for computing a measure of
remaining bandwidth of the communication channel assigned to the community;
and
 item selection apparatus for choosing at least one item to be
30 broadcast from among said plurality of available items based, at least in part, on
said affinity associated with said one item and on said measure of remaining
bandwidth.

18. A system for selecting an item to be broadcast in a broadcasting system, the system for selecting comprising:

community definition apparatus for defining a plurality of
5 communities;

channel assignment apparatus for assigning a communication channel having a bandwidth to each one of said plurality of communities;

affinity apparatus operative, for each one item of a plurality of available items, to determine an affinity between said one item and each one of
10 said plurality of communities, thereby producing a plurality of item-community affinities, and associating said plurality of item-community affinities with said one item; and

community bandwidth apparatus operative, for each one community of said plurality of communities, to compute a measure of remaining bandwidth of
15 the communication channel assigned to the one community and to choose at least one item to be broadcast from among said plurality of available items based, at least in part, on said item-community affinity, associated with said one item, between said at least one item and said one community and on said measure of remaining bandwidth.

20

19. A bandwidth allocation system comprising:

user information apparatus operative to receive user information characterizing preferences of a plurality of broadcast item users;

community information processing apparatus for processing the
25 user information to produce community information, the community information comprising, for each one of a plurality of communities, community profile information;

bandwidth apparatus for assigning a bandwidth to each one of the plurality of communities;

30 item profile apparatus for assigning an item profile to each of a plurality of items to be broadcast; and

item choosing apparatus for choosing at least some of the plurality of items to be broadcast based, at least in part, on a measure of affinity between each of the plurality of items to be broadcast and the community profile information for each one of the plurality of communities, and, at least in part, on the bandwidth assigned to each one of the plurality of communities.

20. Apparatus according to any of claims 17 - 19 and substantially as described hereinabove.

10 21. Apparatus according to any of claims 17 - 19 and substantially as shown in the drawings.

22. A method according to any of claims 1 - 16 and substantially as described hereinabove.

15

23. A method according to any of claims 1 - 16 and substantially as shown in the drawings.

Respectfully submitted,



INVESTOR IN PEOPLE

Application No: GB 0005570.7
Claims searched: all

Examiner: Martyn Dixon
Date of search: 1 November 2000

Patents Act 1977 Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK CI (Ed.R): H4R (RLS,RCSS,RCST,RCT,RCX); H4F (FBB)

Int CI (Ed.7): G06F (17/60); H04H (1/00,1/02); H04N (7/173)

Other: Online: EPODOC,WPI,JAPIO

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	WO 98/37696 A (Herz) see fig 8	
A	WO 98/28869 A (Hyundai)	
A	WO 97/17774 A (Seiko)	
A	US 5758257 A (Herz <i>et al</i>)	

42

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.